

Ifu

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

Date: August 10, 2005

Herkersdorf

Group Art Unit: 2154

Serial No.: 10/658,311

Examiner:

Filed: September 9, 2003

Docket No.: CH20020009US1

For: METHOD FOR PROCESSING A DATA PACKET

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

CERTIFICATE OF MAILING UNDER 37 CFR 1.8 (a)

I hereby certify that the attached correspondence comprising:


Notice Under 37 CFR 1.251
Reply to Notice Under 37 CFR 1.251
Certified copy of European Application No. 02405795.2
Return Postcard

is being deposited with the United States Postal Service as first class mail in an envelope addressed to:

**Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450.**

on August 10, 2005

Allison Berkman
(Type or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

This Page Blank (uspto)



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

Herkersdorf

Serial No.: 10/658,311

Filed: September 9, 2003

For: METHOD FOR PROCESSING A DATA PACKET

Date: August 10, 2005

Group Art Unit: 2154

Examiner:

Docket No.: CH20020009US1

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REPLY TO NOTICE UNDER CFR 1.251-PENDING APPLICATION

Sir:


In response to the Notice Under 37 CFR 1.251-Pending Application, we are enclosing a certified copy of the priority document for filing for the above-identified application. This copy is a complete and accurate copy of the applicant's record of such paper.

Respectfully submitted,

Louis P. Herzberg
Reg. No. 41,500
Tel. (845) 352-3194

IBM CORPORATION
Intellectual Property Law Dept.
P.O. Box 218
Yorktown Heights, New York 10598
Tel: 914 945-2885

This Page Blank (uspto)



COMMISSIONER FOR PATENTS

Approved for use through 07/31/2006. OMB 0651-0031
U.S. Patent and Trademark Office; U. S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Re Application of: HERKERSDORF, ANDREAS et al

Application No. 10658311

Filing Date: 09/09/2003

Title: METHOD FOR PROCESSING A DATA PACKET

Direct to: U.S. Patent and Trademark Office
Mail Stop: **Duplicates**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

NOTICE UNDER 37 CFR 1.251 – Pending Application

Statement (check the appropriate box):

☐ The copy submitted with this reply is a complete and accurate copy of applicant's record of all of the correspondence between the Office and the applicant for the above-identified application (except for U.S. patent documents), and applicant is not aware of any correspondence between the Office and applicant for the above-identified application that is not among applicant's records.

☒ The copy of the paper(s) listed in the notice under 37 CFR 1.251 is/are a complete and accurate copy of applicant's record of such paper(s).

☐ The papers produced by applicant are applicant's complete record of all of the correspondence between the Office and the applicant for the above-identified application (except for U.S. patent documents), and applicant is not aware of any correspondence between the Office and the applicant for the above-identified application that is not among applicant's records.

☐ Applicant does not possess any record of the correspondence between the Office and the applicant for the above-identified application.

Done Aug 9, 05
Date

Louis Herzberg
Signature
Louis P. Herzberg
Typed or printed name

A Copy of this notice should be returned with the reply.

Burden Hour Statement: This collection of information is required by 37 CFR 1.251. The information is used by the public to reply to a request for copies of correspondence between the applicant and the USPTO in order to reconstruct an application file. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This form is estimated to take 60 minutes to complete. This time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450.

PTO Doc Code: R251.RES (Rev. 03/05)



C# 920020003



**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

02405795.2

**CERTIFIED COPY OF
PRIORITY DOCUMENT**

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

This Page Blank (uspto)



Anmeldung Nr:
Application no.: 02405795.2
Demande no:

Anmeldetag:
Date of filing: 12.09.02
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
ETATS-UNIS D'AMERIQUE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se référer à la description.)

Method for processing a data packet

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s)
revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

H04L12/56

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR

This Page Blank (uspto)

METHOD FOR PROCESSING A DATA PACKET

The invention relates to a method for processing a data packet towards a routing destination. More specifically the invention relates to a method for processing a data packet towards a routing destination, wherein a default-route-prefix is determined for specific data packets.

5 TECHNICAL FIELD AND BACKGROUND OF THE INVENTION

10 IP-Forwarding is the process in which an Internet router uses the IP destination address of an incoming packet to search a routing-table using a longest-matching prefix search method in order to determine the next router, i.e. next hop, to which the packet shall be forwarded. If no matching prefix is found then the packet is routed according to a so-called "default route".

15 Two developments can be observed in the routing-table lookup function in Internet routers: First, high search rates for increasing link speeds (OC-192, OC-768, etc.) force the lookup function to use fast but expensive memory such as embedded DRAM and SRAM, having only limited storage capacity. Second, exponentially growing routing-tables force the lookup function to support very large routing-tables. These two developments drive a renewed interest in the use of caches for routing-table lookups.

EP 1122927 is directed to a route lookup engine for determining a next hop. That route lookup engine receives a lookup key and performs a multi-bit trie search with prefix expansion and capture of a variable stride trie.

20 EP 1011231 describes a method an apparatus providing for router redundancy of non internet protocols using the virtual router redundancy protocol.

25 US20020091856A1 describes a default route coding method. A multilevel lookup table includes a plurality of search levels with each search level including a plurality of subtrees, each subtree representing a plurality of nodes. A search of the multilevel lookup table for an entry corresponding to a search key results in a value stored in an entry associated with the node in a subtree. A default value is associated with the root of the subtree. Multiple entries for the subtree can store the default value. To minimize route update time, the default value associated with the subtree is stored in a single location. Instead of storing the default value in multiple entries, each entry stores a use default indicator to indicate that the default value

stored in the single location is to be used. To further reduce the number of locations to modify to update the default route, the single location can store an inherit indicator to indicate that the default value for the subtree is inherited from another subtree.

5 Within enterprise networks, a significant portion of the traffic is likely to be routed according to the default route. The default route, however, does not relate to a specific prefix and is therefore not cached. Consequently, packets that are routed according to the default route, will always result in cache misses.

OBJECT AND ADVANTAGES OF THE INVENTION

10 According to a first aspect of the invention as set forth in the claims a method is provided that provides a higher cache hit rate for lookup operations on a routing table and its routing table cache. This is effected by the fact that the method provides default-route prefixes for default-route-directed data packets. The default-route prefixes can be fed to the cache such that subsequent default-route-directed data packets have an increased likelihood of encountering a cache hit and being directed to the default route without having to perform a
15 lookup operation on the routing table. In the following, a matching prefix is defined as a prefix that is a prefix to the destination address.

It is of advantage if in a first lookup step for the destination address the matching destination address prefix is searched in the routing table cache, and if said first lookup step 1 results in not finding such matching destination address prefix, in a second lookup step for the
20 destination address the matching destination address prefix is searched in the routing table, because by this two-step lookup process, there will be fewer second lookup steps counted over a multitude of data packets which results in a quicker and less resource-intensive forwarding process.

It provides advantageous if the second lookup step does not result in a matching destination
25 address prefix, to forward in a default forwarding step the data packet to a default routing destination, because it has been recognized that there is no matching destination address prefix and it makes more sense to forward the data packet to a default destination for further processing, than deleting it. The default destination need not be identical for different data packets. The system may use various default destinations and a selection method may be

applied to determine which of the default destinations shall be applied to a specific data packet.

If in a default-route caching step, the default-route-prefix is entered together with the default routing destination as an entry into the routing table cache, the process for forwarding data packets is accelerated, since subsequent data packets heading for a destination address that was not contained in the routing table or its cache, but is covered by the entered default-route-prefix will now encounter a cache hit due to the default routing destination entry in the routing table cache, and hence be forwarded to the default destination without having to perform a second lookup step on the routing table. If, however, such a subsequent data packet with a different destination address lacks a matching destination address prefix in the routing table cache and in the routing table, this meaning that its destination address is not covered by the entered default-route-prefix, then another default-route-prefix is determined and entered together with the default routing destination as an entry into the routing table cache. The routing table cache may hence contain several default-route-prefix entries. These need not necessarily have the same default routing destination.

If the second lookup step on the routing table results in finding a matching destination address prefix, the found entry is entered into the routing table cache in a cache update step, and the data packet is forwarded in a destination forwarding step to the corresponding routing destination. This procedure provides the advantage that so-called locality is exploited, i.e., it is likely that several data packets in a short time-distance to each other. Such data packets will encounter cache hits and be forwarded to their routing destination without having to access the routing table. This saves time and resources.

If in the first lookup step the routing table cache is searched also for covering path entries that reside in the routing table cache, in their totality covering at least all destination address prefixes existing in the routing table, the cache-misses will be usable for direct default-path forwarding because the covering path entries provide the assurance that in the event of a cache miss, the searched destination address will not exist in the routing table. Hence, a direct default-path forwarding can follow.

In the event that the first lookup step results in finding a matching destination address prefix, the data packet can be forwarded in a destination forwarding step to the

corresponding routing destination. This saves time and resources since no second lookup step is necessary.

If in the event that the first lookup step results in finding a covering path entry for the destination address, in a second lookup step for the destination address the matching destination address prefix is searched in the routing table, this second lookup step provides the final distinction between those data packets that encounter a valid routing destination and those data packets which shall be forwarded to the default routing destination.

SUMMARY OF THE INVENTION

The invention is directed to a method for processing a data packet that has a destination address. In the event the destination address lacks a matching destination address prefix in a routing table cache and in a routing table, a default-route-prefix is determined in a default-route determination step. The invention is furthermore directed towards a computer program product comprising program code means for performing the method and a computer program product comprising the program code means stored on a computer-readable medium.

DESCRIPTION OF THE DRAWINGS

Examples of the invention are depicted in the drawings and described in detail below by way of example. It is shown in

Fig. 1 a flow diagram of a method for accelerating default-route packet-forwarding,

Fig. 2 a block diagram of an arrangement for routing data packets in accordance with a routing table,

Fig. 3 a tree diagram of an address space covered by an exemplary set of entries in a routing table,

Fig. 4 the tree diagram of fig. 3 with a default path entry,

Fig. 5 a block diagram of a multi-bit address space covered by an address,

Fig. 6 a tree diagram of the address space pertaining to fig. 5,

Fig. 7 a flow diagram of an alternative method for accelerating default-route packet-forwarding,

Fig. 8 a tree diagram of an address space covered by an exemplary set of entries in a routing table with cache entries in accordance with the method of fig. 7.

- 5 All the figures are for sake of clarity not shown in real dimensions, nor are the relations between the dimensions shown in a realistic scale.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In the following, the various exemplary embodiments of the invention are described.

- 10 In figure 1 a flow diagram of a method of accelerating default-route packet-forwarding is depicted.

- 15 A data packet heading for a destination is expected to arrive at an input port of a router. It is the task of the router to direct the arrived packet in accordance with a destination address that the data packet carries with it, to a corresponding routing destination. To achieve this, the router comprises a routing table with entries comprising destination address prefix - routing destination pairs.

- 20 In the event that a data packet arrives with a destination address, a first lookup step 1 is performed during which for the destination address d a matching destination address prefix is searched in a routing table cache L_1 . A match is defined as the case when the destination address prefix is identical to the corresponding number of bits of the destination address d . In many cases the destination address prefix will be shorter than the destination address d .

If the first lookup step on the routing table cache L_1 results in a hit, i.e. finding a matching destination address prefix, labeled with "y" in fig. 1, the data packet is forwarded in a destination forwarding step 7 to the corresponding routing destination.

- 25 If the first lookup step 1 on the routing table cache L_1 results in a miss, labeled with "n" in fig. 1, the method proceeds with a second lookup step 2. In the second lookup step 2, for the destination address d a matching destination address prefix is searched in the routing table L_2 . If the second lookup step 2 on the routing table L_2 results in a hit, i.e. finding a matching destination address prefix, labeled with "y" in fig. 1, the found entry is entered into the

routing table cache L_1 in a cache update step 6, and the data packet is forwarded in the destination forwarding step 7 to the corresponding routing destination.

If the second lookup step 2 does not result in a matching destination address prefix, labeled with "n" in fig. 1, the corresponding data packet will in subsequent steps be routed according to the default route, i.e. to a default routing destination.

In a default-route determination step 3, a default-route-prefix P_d is derived. Hence, in this step, for the destination address d lacking a matching destination address prefix in the tables L_1 , L_2 , the default-route-prefix P_d is determined. The default-route-prefix P_d comprises, as its name states, in most cases only a subset of the number of bits that the destination address d has. It is of advantage to determine the default-route-prefix P_d shorter than the destination address d , such that the default-route-prefix P_d covers more than one destination address, taking into account however, that the default-route-prefix P_d may only cover destination addresses that have no matching destination address prefix in the routing table L_2 . It is therefore the preferred embodiment to determine the default-route-prefix P_d to be the shortest possible default-route-prefix P_d that satisfies the above condition. Thereby the default-route-prefix P_d covers the maximum of destination addresses.

In a default-route caching step 4, the default-route-prefix P_d is entered together with the default routing destination as an entry into the routing table cache L_1 .

In a default forwarding step 5 the data packet is afterwards forwarded to the default routing destination.

In figure 2, a router is depicted with a data packet input 11 leading via a route lookup engine 12 and a packet forwarding engine 13 to a data packet output 14. The route lookup engine 12 has access to the routing table L_2 and the routing table cache L_1 . Once a data packet has arrived via the data packet input 11, the route lookup engine 12 uses the destination address for its lookup operation, accessing the routing table L_2 and the routing table cache L_1 . Dependent on the result of the lookup operation, the route lookup engine 12 delivers a routing destination address for the data packet. The packet forwarding engine 13 uses that routing destination address to direct the data packet to a data packet output 14.

The method of accelerating default-route packet-forwarding will hereinafter be explained in more detail using the example of a longest-matching prefix search, using a 4-bit search key, on the routing table L_2 containing here the following three prefixes in binary notation:

5 prefix 1: 001 (length 3)
 prefix 2: 0111 (length 4)
 prefix 3: 11 (length 2)

Figure 3 shows the corresponding binary tree and the portions of the 4-bit address space AS that are covered by these prefixes. The 4-bit address space AS comprises here six portions I, II, III, IV, V, VI, three of which, I, III, V are not covered by the prefixes.

10 These uncovered portions I, III, V of the address space AS correspond to the default route. The default route is the route that is selected when no matching prefix is found in the routing table L_2 .

15 The uncovered address space comprises here the addresses “0000”, “0001”, “0100” to “0110”, and “1000” to “1011”. These addresses and address intervals can also be represented by a set of prefixes, being referred to as default-route-prefixes. If all those default-route-prefixes would be added to the routing table L_2 then always a matching prefix will be found during the second lookup step 2. This would eliminate the notion of a default-route corresponding to the situation that no matching prefix is found. However, the total number of those default-route-prefixes very likely outnumber the number of the
 20 non-default-route prefixes in actual routing tables. It would therefore be impractical to add all default-route-prefixes to the routing table L_2 because this would significantly increase the required storage capacity of the routing table L_2 and degrade the update performance.

25 Instead, the following approach is taken here. The data packet arrives with its destination address. The first lookup step 1 is performed on the routing table cache L_1 . During the first lookup step 1, for the destination address a matching prefix is searched in the routing table cache L_1 . Assuming that the first lookup step 1 on the routing table cache L_1 results in a miss, i.e. no matching prefix is found, the second lookup step 2 is performed. The second lookup step 2 is executed on the routing table L_2 . Assuming that the lookup step 2 on the routing table L_2 does not result in a matching prefix, this means that the arrived data packet

will be routed according to the default route. In the default-route determination step 3 a default-route-prefix P_d is derived. Hence, in this step, for the destination address lacking a matching prefix, the default-route-prefix P_d is determined.

5 Figure 4 shows an example of how the default-route-prefix P_d can be derived on-the-fly during the lookup step 2 on the binary tree shown in figure 3. This example involves the use of a search key "1010", i.e. corresponding to the situation that the data packet has arrived with the destination address "1010". The lookup step 2 will stop in the tree node that is pointed to in fig. 1 by the large arrow, because of the second search key bit being a "0". Thereafter, all the search key bits up to and including the bit on which the lookup step 2 was
10 stopped because it did not match with a subsequent child node, are determined to be the default-route-prefix P_d . This default-route-prefix P_d is the shortest possible default-route-prefix P_d that covers only destination addresses that have no matching prefix in the routing table L_2 .

15 In the default-route caching step 4, the default-route-prefix P_d is entered as an entry into the routing table cache L_1 , together with the default routing destination. Storing the default-route-prefix P_d , that covers more than one destination address, in the routing table cache L_1 , saves storage space since for several destination addresses only one entry resides in the routing table cache L_1 .

20 The default-route-prefix P_d equals "10" in the example given in fig. 4, covering four destination addresses. In fig. 4 this default-route-prefix P_d is shown together with the portion of the address space that it covers. In this example, the address space covered by the default-route-prefix P_d equals the uncovered portion V of the address space AS .

25 The default-route-prefix P_d will only be entered into the routing table cache L_1 and not be stored as an actual prefix in the routing table L_2 . For later arriving data packets, the search keys that fall into the portion V of the address space AS covered by this particular default-route-prefix P_d will experience a cache hit as long as this default-route-prefix P_d stays cached. Such search keys are likely to occur when so called locality of reference exists, which is a main motivation for using a cache.

30 This means that for the next data packet, in the event that the destination address is not contained in the routing table L_2 , but in the event that the destination address is nevertheless

contained in the address space covered by the default-route-prefix P_d whose entry is in the routing table cache L_1 , the first lookup step 1 will result in a cache hit and lead to the destination forwarding step 7 that directly forwards the data packet to the default routing destination, without having had to execute the second lookup step 2. Each

5 default-route-prefix entry in the routing table cache L_1 , can in principle have a different default routing destination. For such case, it may also be of advantage not to store one entry for the shortest-possible default-route-prefix P_d in the routing table cache L_1 , but to split the address space it covers up into several address spaces covered by several default-route-prefix entries with different destination addresses.

10 The default-route-prefix P_d can be derived in a similar manner in an arrangement of a multi-bit trie. Figure 5 shows an example of such a trie, implemented as a data structure in which search key segments of an IPv4 destination address are used to index tables that are linked together to form a hierarchical structure. The IPv4 destination address comprises a first segment S_1 , a second segment S_2 , and a third segment S_3 . The first segment S_1 covers

15 the address space of a first table T_1 , the second segment S_2 covers the address space of a second table T_2 , and the third segment S_3 covers the address space of a third table T_3 .

The data structure in Figure 5 contains the following prefixes in hexadecimal notation:

prefix	prefix length	search result
123456	24	P
1234567	28	Q
123456AB	32	R

Figure 6 illustrates the portions of the 32-bit address space that are covered by these prefixes. Taking an example of a lookup on an IPv4 destination address "123489AB", the

25 lookup step 2 will not find a matching prefix in this data structure. This lookup step 2 will stop in the table that is indexed by the second IPv4 destination address segment S_2 , since this table does not contain a valid entry at the offset "89". The only valid entry is contained at the offset "56".

The default-route-prefix is here determined by taking all the IPv4 address segments up to the one that caused the lookup step 2 to stop. In this example, the default-route-prefix would consist of the first segment S_1 and the second segment S_2 of the IPv4 destination address, resulting in a default-route-prefix "123489". The portion of the address space covered by this default-route-prefix is illustrated in figure 6 and ranges from "12348900" to "123489FF".

The above method hence can be summarized in that a default-path entry is created in the routing table cache L_1 upon a data packet having been recognized as having to be forwarded along the default-path. The default path entry is advantageously the shortest possible path covering only destination addresses that pertain to the default path. Thereby, the processing time for subsequent data packets with their destination address being covered by the default path is reduced.

In fig. 7 a flow diagram of the steps of a method for accelerating default-route packet-forwarding is shown. As far as identical steps are performed, the reference numbers of fig. 1 have been maintained. In the first lookup step 1 the destination address of the data packet is used for searching a thereto-matching destination address prefix in the routing table cache L_1 . For the searching there exist three possible results: A cache hit can occur, i.e. a valid destination address prefix can be found, designated with y_1 in fig. 7. In this case, the data packet will be forwarded in the destination forwarding step 7 to the corresponding destination address. As a second possible result, a cache miss can be the case, designated with n in fig. 7. In that case, the routing table L_2 need not be searched and the data packet is forwarded to the default routing destination in the default forwarding step 5. As a third possible search result, a covering path entry can be found, designated with y_2 in fig. 7. In that case, the routing table L_2 need be searched in order to know whether within the coverage of the covering path entry there is a valid destination address prefix for the destination address or not. Hence in this case the second lookup step 2 is performed. If that step results in a hit, designated with y in fig. 7, the found destination address prefix is entered into the routing table cache L_1 in the cache update step 6 and the data packet will be forwarded to the corresponding destination address in the destination forwarding step 7. If the second lookup step 2 results in a miss, designated with n in fig. 7, the data packet is forwarded to the default routing destination in the default forwarding step 5.

The routing table cache L_1 stores in this example not only the entries of destination addresses but also additional entries that cover all prefixes that exist in the routing table L_2 , herein referred to as covering path entries. When a covering path entry is found as a result of the first lookup step 1, i.e. being a cache hit, only then will a lookup operation on the routing table L_2 be performed. This lookup can still result in the default route. However, with this method a cache miss relates always to the default route.

In fig. 8 an example is illustrated giving an address space for a four-bit tree. The covering path prefixes stored in the routing table cache L_1 to cover all valid prefixes are designated with y_2 . The covering path prefixes are stored in the routing table cache L_1 at the time when the routing table is loaded with valid destination address prefixes. Updating of the covering path prefixes is only necessary when the routing table L_2 content is changed. The length of the covering path prefixes can be selected. The longer the covering path prefix, the more such covering path prefixes will be necessary to cover all valid destination address prefixes in the routing table L_2 and on the other hand the lower is the probability of the y_2 case, which is the case when the second lookup step 2 is performed. The selection of the length of the covering path prefix is hence a tradeoff between cache size to maintain the covering path prefixes and processing time for the lookup operation. It is seen as an advisable solution to select the length of the covering path prefixes such that the ratio between the valid destination address prefixes and the default routing destinations covered by that covering path prefix is at least 1:2.

Comparing this method with the method depicted in fig. 1, it becomes apparent that the result y_1 in the first lookup step 1 in fig. 7 is equivalent to the result y in the first lookup step 1. The result n in the first lookup step 1 in fig. 7 is equivalent to the result n in the first lookup step 1 combined with the result n in the second lookup step 2. It is possible to add to the method of fig. 7 the default-route determination step 3 and the default-route caching step 4. The result y_2 in the first lookup step 1 in fig. 7 is equivalent to the result n in the first lookup step 1, searching in the second lookup step 2, resulting in y , and arriving at the cache update step 6.

This method breaks with the traditional notion of a cache, as the routing table cache L_1 is now used as an isolated buffer memory, and the routing table L_2 is only accessed when this

is directed specifically by the output of the cache lookup. This method is combinable with the method described in conjunction with fig. 1.

Another possibility to enter the default-route-prefix P_d into the routing table cache L_1 is to not wait until a data packet arrives, but to perform an upfront check on the routing table L_2 in which the address space of the routing table L_2 is analyzed for valid destination address prefixes, and uncovered prefixes are used for determining default-route-prefixes P_d therefor and creating an entry comprising the default-route-prefix P_d together with a default routing destination, in the routing table cache L_1 . This is also referred to as prefetching method and it has the advantage of enabling faster data packet processing. It may even be combined with the data-packet triggered determination of the default-route-prefix P_d , in that only a part of the uncovered prefixes is subjected to the prefetching method and the rest is subject to the data-packet triggered determination. Prefetching provides more advantageous for the parts of the address space that can be covered by a short default-route-prefix P_d , since that prefix uses reduced table space. Having used the prefetching method, a default-route-prefix P_d already resides together with a default routing destination as an entry in the routing table cache L_1 , at the time of arrival of the data packet and in the event that the default-route-prefix P_d matches with at least part of the destination address d , in the default forwarding step 5 the data packet is forwarded to the corresponding default routing destination.

It is obvious for the person skilled in the art that the present invention can be realized in hardware, software, or a combination of these. Also, it can be implemented in a centralized fashion on one single computer system, or in a distributed fashion where different elements are spread across several interconnected computers or computer systems, whereby any kind of a computer system - or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods.

Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b)
5 reproduction in a different material form.

Any disclosed embodiment may be combined with one or several of the other embodiments shown and/or described. This is also possible for one or more features of the embodiments.

It is obvious that a person skilled in the art can modify the shown arrangements in many ways without departing from the gist of the invention which is encompassed by the
10 subsequent claims.

This Page Blank (uspto)

CLAIMS

1. Method for processing a data packet, having a destination address (d), towards a routing destination, wherein in the event that in a routing table cache (L_1) and in a routing table (L_2), there is no entry with a destination address prefix that is a prefix of the destination address (d), a default-route-prefix (P_d) is determined in a default-route determination step (3).
5
2. Method according to claim 1, wherein the default-route-prefix (P_d) is determined be a prefix of at least the destination address (d), preferably being also a prefix of other destination addresses for which in the routing table (L_2) there is no entry with a destination address prefix that is a prefix of the other destination address.
10
3. Method according to claim 1 or 2, wherein in a first lookup step (1) for the destination address (d) the destination address prefix being a prefix thereof is searched in the routing table cache (L_1), and wherein if said first lookup step (1) results in not finding such destination address prefix, in a second lookup step (2) for said destination address (d) the destination address prefix being a prefix thereof is searched in the routing table (L_2).
15
4. Method according to claim 3, wherein if the second lookup step (2) on the routing table (L_2) results in finding the destination address prefix being a prefix of the destination address (d) a matching destination address prefix, the found destination address prefix entry is entered into the routing table cache (L_1) in a cache update step (6), and the data packet is forwarded in a destination forwarding step (7) to the corresponding routing destination.
20
5. Method according to claim 3 or 4, wherein if the second lookup step (2) results in not finding the destination address prefix being a prefix of the destination address (d), in a default forwarding step (5) the data packet is forwarded to a default routing destination.
25

6. Method according to one of claims 1 to 5, wherein in a default-route caching step (4), the default-route-prefix (P_d) is entered together with the default routing destination as an entry into the routing table cache (L_1).
- 5 7. Method according to claim 3, wherein in the first lookup step (1) the routing table cache (L_1) is searched for covering path entries that reside in the routing table cache (L_1), the covering path entries in their totality being a prefix for at least all destination address prefixes existing in the routing table (L_2).
- 10 8. Method according to claim 7, wherein in the event that the first lookup step (1) results in finding no covering path entry for the destination address (d), the data packet is forwarded to a default routing destination in a default forwarding step (5).
9. Method according to claim 7 or 8, wherein in the event that the first lookup step (1) results in finding a covering path entry for the destination address (d), in a second lookup step (2) for said destination address (d) the destination address prefix being a prefix of the destination address (d) is searched in the routing table (L_2).
- 15 10. Method according to one of claims 3 to 9, wherein in the event that the first lookup step (1) results in finding the destination address prefix being a prefix of the destination address (d), the data packet is forwarded in a destination forwarding step (7) to the corresponding routing destination.
- 20 11. Method for processing a data packet, having a destination address (d), towards a routing destination, wherein a default-route-prefix (P_d) resides together with a default routing destination as an entry in a routing table cache (L_1) and wherein in the event that the default-route-prefix (P_d) matches with at least part of said destination address (d), in a default forwarding step (5) the data packet is forwarded to said default routing destination.

12. Computer program product comprising program code means for performing a method according to one of the preceding claims.
13. Computer program product according to claim 12, comprising the program code means stored on a computer-readable medium.

ABSTRACT

The invention is directed to a method for processing a data packet that has a destination address. In the event that in a routing table cache and in a routing table, there is no entry with a destination address prefix that is a prefix of the destination address, a default-route-prefix is determined in a default-route determination step.

5

This Page Blank (uspto)




Fig. 1






Fig. 2






Fig. 5




Fig. 6




Fig. 7




Fig. 8